

Holographic Quantum Error Correction on Bruhat–Tits Trees: A Conditional Threshold and Resource Analysis at 10^{-4}

Holographic Quantum Error Correction on Bruhat–Tits Trees: A Conditional Threshold and Resource Analysis at 10^{-4}

Abstract

We assess, quantitatively and conditionally, whether a quantum error correction (QEC) architecture in which the encoding map is realized as renormalization-group (RG) flow on a Bruhat–Tits tree $\mathcal{T}_p$ — the p -adic analogue of hyperbolic anti-de Sitter space — can achieve a QEC threshold of at least 10^{-4} . Because no hardware data or full error-model simulation is supplied with the proposal, we construct an explicit conditional model: independent site errors on a tree of valency $q = p + 1$, a concatenation-style logical error law $\epsilon_L = A(\epsilon/\epsilon_{th})^{(d+1)/2}$, and a target failure probability of 0.01 over 10^9 correction cycles. Under these assumptions we derive, with full arithmetic, that (i) at a physical error rate exactly equal to 10^{-4} the model provides no protection, so the threshold claim is meaningful only strictly below 10^{-4} ; (ii) at $\epsilon = 10^{-5}$ a code distance $d = 21$ suffices, requiring 88,573 tree sites per logical qubit on $\mathcal{T}_3$ and a measurement bandwidth of 8.8573×10^{10} events per second per logical qubit at a 10^6 Hz cycle rate; (iii) a shallower $d = 11$ code meets the same target only if $\epsilon \leq 1.4678 \times 10^{-6}$. A complementary convention modeling the RG step as a quadratic map $p_{\ell+1} = \kappa p_\ell^2$ yields the equivalent condition that the threshold equals $1/\kappa$, so the 10^{-4} target holds if and only if $\kappa \leq 10^4$. We conclude that the threshold hypothesis is internally consistent but operationally demanding, and we identify the assumptions whose failure would falsify it.

1. Introduction

The question we address is quantitative: does the holographic QEC architecture of [9] — in which quantum error correction is organized as renormalization-group (RG) flow on a Bruhat–Tits tree $\mathcal{T}_p$, the p -adic analogue of hyperbolic anti-de Sitter (AdS) space, so that the encoding map of a holographic code is literally an RG trajectory from boundary degrees of freedom into the bulk — achieve a QEC threshold of at least 10^{-4} ?

A threshold, in the QEC sense used throughout this paper, is a physical error rate ϵ_{th} below which increasing code distance d drives the logical error rate ϵ_L to zero. The claim under test is that $\epsilon_{th} \geq 10^{-4}$ for this architecture. This is a claim about an error-correcting *architecture*, and it must be answered with the same discipline that the systems-architecture literature applies to other resource-constrained designs: explicit overhead accounting, explicit bandwidth accounting, and explicit statement of the regime in which the design works. The resource-burden perspective of [10] makes this discipline concrete: active QEC protects quantum information only by continuously consuming physical resources — qubits, measurement bandwidth, classical processing, and energy — so any threshold claim must be paired with the cost of operating at that threshold. The same accounting imperative appears in the quantum-architecture literature: the main obstacle to practical quantum computing is limited qubit count, which motivates floorplans that minimize memory overhead while preserving fault tolerance [5].

Because the proposal is new and no experimental or simulated error data is supplied with it, every number in this paper is either (a) an input stated explicitly as an assumption of the conditional model, or (b) derived from such inputs by arithmetic shown in full in Section 4. We do not claim to have measured or simulated anything. Our contribution is a transparent conditional analysis: given the model of Section 3, the threshold hypothesis at 10^{-4} has precise, checkable consequences, and we compute them.

The paper proceeds as follows. Section 2 situates the work relative to the supplied literature. Section 3 defines the conditional model and a complementary alternative convention. Section 4 carries out all derivations with explicit arithmetic. Section 5 reports only the computed numbers. Section 6 discusses limitations, failure modes, and falsification conditions. Appendix A documents points where independent drafts of this analysis diverged and the conventions adopted here; Appendix B attributes each substantive claim to its sources.

2. Background and Related Work

We discuss all twelve supplied works. Because several entries carry only truncated summaries, we attribute to each work only what its own supplied text states, and say so explicitly where the summary is thin.

[9] Holographic QEC as AdS/CFT RG flow on Bruhat–Tits trees. This is the proposal under assessment. Its supplied summary states that QEC is organized as RG flow on a Bruhat–Tits tree $\mathcal{T}_p$, the p -adic analogue of hyperbolic AdS space, so that the encoding map of a holographic code is literally an RG trajectory from boundary degrees of freedom into the bulk, and that the work performs a *conditional* threshold analysis at 10^{-4} . The summary is truncated mid-sentence and supplies no further detail on the error model or the threshold derivation; our paper is therefore an independent, fully explicit reconstruction of what such a conditional analysis must contain, anchored to the framing that [9] states.

[10] The resource burden of active QEC. The supplied summary states that active QEC protects quantum information only by continuously consuming physical resources — qubits, measurement bandwidth, classical processing, and energy — and that while the theoretical foundations of QEC are mature, the *resource burden* of running a code, rather than constructing one, is the object of study. This is exactly the accounting frame we adopt in Sections 4 and 5: a threshold number is meaningless without the qubit count and measurement bandwidth needed to operate below it. The summary is truncated and gives no specific overhead formulas, so we derive our own.

[11] Coherence-assisted error correction for deep learning. The supplied summary describes a hybrid quantum-classical architecture in which deep learning inference and training are limited by the energy cost of protecting computation against bit-level faults from aggressive low-voltage operation, near-threshold computing, or radiation-prone deployment, and states that the paper reconciles a feasibility analysis (text truncated). Its relevance is methodological: it treats error correction as an architecture-level co-design problem with an explicit feasibility verdict, the posture we take toward the 10^{-4} question. The truncated summary supplies no quantitative results we could cite.

[12] **Lifecycle of a fault-tolerant quantum computer.** The supplied entry consists of a title only; its summary field is empty. We therefore cannot state anything about its content beyond what the title asserts. We cite it to acknowledge that fault-tolerant systems plausibly have a full operational lifecycle, and note that our analysis covers only the steady-state operation segment. No claim about this work's findings is made or implied.

[5] **LSQCA: resource-efficient load/store architecture for limited-scale fault-tolerant quantum computing.** The supplied summary states that current fault-tolerant quantum computer (FTQC) architectures use several encoding techniques to enable reliable logical operations with restricted qubit connectivity, that such logical operations demand additional memory overhead to ensure fault tolerance, and that because the main obstacle to practical quantum computing is limited qubit count, the work's primary mission is to design floorplans that reduce that overhead (text truncated). This is the closest quantum-architectural neighbor of our problem: it treats fault tolerance as a floorplan and memory-overhead problem, exactly as we treat the Bruhat–Tits tree volume as a qubit-overhead problem in Section 4. The truncated summary gives no specific floorplan numbers.

[8] **Insights into DeepSeek-V3.** The supplied summary states that rapid scaling of large language models has unveiled critical limitations in current hardware architectures, including constraints in memory capacity, computational efficiency, and interconnection bandwidth, and that DeepSeek-V3, trained on 2,048 NVIDIA H800 GPUs, demonstrates how hardware-aware model co-design can address these challenges, enabling cost-efficient training and inference (text truncated). We cite it as a modern example of co-design reasoning at scale: the lesson that architecture and workload must be designed jointly transfers directly to QEC, where the "workload" is the error channel and the "hardware" is the tree interconnect. No specific DeepSeek-V3 performance figures are used here.

[2] **Architectural implications of graph neural networks.** The supplied summary states that graph neural networks (GNNs) operate on graph structures, are popular due to high accuracy on graph-related tasks, but are not as well understood in the systems and architecture community as multi-layer perceptrons and convolutional neural networks, and that the work tries to improve that understanding (text truncated). Our architecture is graph-structured (a tree), and any future classical co-processor that decodes it would face exactly the GNN-style question of how graph-structured workloads stress memory hierarchies. The truncated summary gives no specific findings.

[3] **Dalorex: data-local program execution for memory-bound applications.** The supplied summary states that applications with low data reuse and frequent irregular memory accesses, such as graph or sparse linear algebra workloads, fail to scale well due to memory bottlenecks and poor core utilization; that prefetching, decoupling, and pipelining can mitigate memory latency but memory bottlenecks persist due to limited off-chip bandwidth (text truncated). Tree-structured QEC decoding is precisely such a workload: syndrome data on a $\mathcal{T}_3$ tree has low reuse and irregular access patterns, so the bandwidth wall [3] identifies is the classical-side analogue of the measurement-bandwidth wall we compute in Section 4.

[6] **Assessing spread in sets of software architecture designs.** The supplied summary states that several approaches generate architecture design alternatives by optimization, aiming to improve an initial architecture with respect to quality aspects such as performance, reliability, or maintainability, and that each optimization experiment usually produces a different set of alternatives (text truncated). We cite it for the meta-level point that a single architecture proposal should be evaluated as one point in a family of design alternatives — here the family $\{\mathcal{T}_p, d, \epsilon\}$, or equivalently the range of the amplification constant κ in Section 4.8 — rather than in isolation. The truncated summary supplies no specific metrics.

[1] **Stochastic model-based proxy servers for VoD.** The supplied summary states that in a video-on-demand system the main repository may be far from the user with limited streaming capacity, that high-quality video requires high bandwidth, and that a distributed architecture with multiple local servers can achieve a higher video hit ratio and reduced client waiting time (text truncated). The structural analogy is direct: a central repository with limited streaming capacity maps onto a central decoder fed by distributed syndrome streams, and the distributed-server remedy maps onto distributing decoding across tree subtrees. The truncated summary gives no quantitative hit-ratio figures.

[7] **Reconfigurable DSP architecture using field-programmable DSP arrays.** The supplied summary states that modern integrated-circuit technology makes cheaper, faster, smaller special-purpose signal-processing circuits feasible, that DSP functions are implemented either on inflexible ASICs or on FPGAs with lower utilization or speed, and that a Field Programmable DSP Array (FPDA) is proposed (text truncated). Real-time syndrome decoding at the bandwidths we compute in Section 4 is a DSP throughput problem, so reconfigurable DSP fabrics are a plausible implementation substrate; the truncated summary gives no utilization numbers.

[4] **QVLM architecture and SQuID dataset.** The supplied summary states that current vision-language models lose pixel-level precision through architectural compression, degrading quantitative spatial reasoning (text truncated). The analogy is cautionary: an architecture that coarse-grains its input — as RG flow on a tree coarse-grains error information — necessarily loses fine-grained detail, and the design question is whether the information needed for the downstream task survives the coarse-graining. Our threshold analysis is exactly a test of what survives the RG coarse-graining of [9]. The truncated summary supplies no further specifics.

In summary: [9] supplies the proposal and the threshold question; [10], [5], and [12] supply the resource- and lifecycle-accounting perspective; [11], [8], [2], [3], [1], [7], [4], and [6] supply architectural analogies from co-design, graph workloads, memory-bound execution, distributed serving, reconfigurable DSP, information loss under compression, and design-space evaluation. No supplied entry provides experimental QEC data; the quantitative content of this paper is therefore entirely our own conditional derivation.

3. Methods

3.1 Convention 1 (primary): distance-based conditional model

Following the framing of [9] — encoding as RG trajectory on $\mathcal{T}_p$ — we model the code as a ball of radius R in the Bruhat–Tits tree $\mathcal{T}_q$ of valency $q = p + 1$, with code distance $d = 2R + 1$. The model assumptions, each stated explicitly, are:

- **(A1) Error channel.** Each bulk site fails independently per correction cycle with physical error rate ϵ (depolarizing-style, symmetric).
- **(A2) Logical error law.** The logical error per cycle follows the standard concatenation-style scaling

$$\epsilon_L(d, \epsilon) = A \left(\frac{\epsilon}{\epsilon_{th}} \right)^{(d+1)/2},$$

where $A \geq 1$ is an architecture-dependent prefactor and ϵ_{th} is the threshold. We take $A = 1$ (the most favorable case) unless stated.

- **(A3) Hypothesis under test.** $\epsilon_{th} \geq 10^{-4}$; we evaluate consequences at the boundary value $\epsilon_{th}^{(0)} = 10^{-4}$.
- **(A4) Reliability target.** Total failure probability over a run of T cycles must satisfy $P_{fail} = 1 - (1 - \epsilon_L)^T \leq 0.01$.
- **(A5) Run length.** $T = 10^9$ correction cycles.
- **(A6) Cycle rate.** $f_c = 10^6$ correction cycles per second (one complete syndrome-extraction round per cycle).
- **(A7) Tree valency.** $p = 2$, so $q = p + 1 = 3$.

Every primary result in Section 5 is conditional on (A1)–(A7). Violating any of them changes the numbers; Section 6 says how.

3.2 Tree volume (qubit overhead)

The number of sites in a ball of radius R in $\mathcal{T}_q$ is the sum of a geometric series:

$$V_q(R) = 1 + q \sum_{j=0}^{R-1} q^j = 1 + q \frac{q^R - 1}{q - 1}.$$

We take $N_{phys} = V_q(R)$ as the physical-site count per logical qubit, in the spirit of the overhead accounting of [10].

3.3 Measurement bandwidth

Each site is measured once per cycle, so the per-logical-qubit measurement event rate is

$$B = V_q(R) f_c.$$

3.4 Distance requirement

From (A2) and (A4), using $1 - (1 - \epsilon_L)^T \approx T \epsilon_L$ for $\epsilon_L T \ll 1$, the requirement is

$$T A \left(\frac{\epsilon}{\epsilon_{th}} \right)^{(d+1)/2} \leq 0.01,$$

so, solving for d given ϵ , ϵ_{th} , A , T :

$$d \geq 2 \frac{\ln(0.01/(TA))}{\ln(\epsilon/\epsilon_{th})} - 1.$$

3.5 Convention 2 (complementary): quadratic RG map

As a complementary view, documented as a divergent convention in Appendix A, we model the RG step directly. At level ℓ of the tree (level 0 = boundary leaves, increasing toward the bulk root), let p_ℓ denote the probability that the effective encoded information at that level has suffered an uncorrectable error. For any code family in which uncorrectable events at a node require errors on a bounded number of its children, the leading-order recursion is quadratic:

$$p_{\ell+1} = \kappa p_\ell^2,$$

where $\kappa > 0$ is the effective amplification constant, absorbing the number of error patterns that map to an uncorrectable event at the parent. Assumptions: **(A1')** the quadratic map holds with κ independent of ℓ ; **(A2')** errors on distinct nodes at the same level are independent; **(A3')** the tree is regular and the code self-similar under the RG step. The threshold is the unstable fixed point of this map, computed in Section 4.8.

4. Analysis

All inputs are stated with their provenance. Inputs from the model assumptions of Section 3 are labeled (A1)–(A7) or (A1')–(A3'); no number comes from any external measurement.

Input 1. $\epsilon_{th}^{(0)} = 10^{-4}$ — hypothesis under test, per [9]'s stated conditional threshold analysis at 10^{-4} (A3).

Input 2. $\epsilon = 10^{-5}$ — assumed operating point, one order of magnitude below the hypothesized threshold (stated assumption, not a measurement).

Input 3. $T = 10^9$ cycles (A5); $P_{fail} \leq 0.01$ (A4); $A = 1$ (A2); $q = 3$ (A7); $f_c = 10^6 \text{ s}^{-1}$ (A6).

4.1 Derivation 1: no protection exactly at the threshold

Set $\epsilon = \epsilon_{th}^{(0)} = 10^{-4}$. Then

$$\frac{\epsilon}{\epsilon_{th}} = \frac{10^{-4}}{10^{-4}} = 1,$$

so for every d ,

$$\epsilon_L = A \cdot 1^{(d+1)/2} = A = 1.$$

A logical error rate of 1 per cycle means the encoded information is destroyed every cycle regardless of distance. **Conclusion:** at $\epsilon = 10^{-4}$ exactly, the architecture provides zero protection; the threshold claim is meaningful only for ϵ strictly below 10^{-4} , with a margin quantified in Derivation 3. (Under Convention 2 the same point is expressed as p^* being an unstable fixed point: at $p_0 = p^*$ the error neither shrinks nor grows. Appendix A records this

interpretive divergence.)

4.2 Derivation 2: required distance at $\epsilon = 10^{-5}$

Requirement: $T \epsilon_L \leq 0.01$, i.e.

$$10^9 \left(\frac{10^{-5}}{10^{-4}} \right)^{(d+1)/2} \leq 10^{-2}.$$

The ratio is $\epsilon/\epsilon_{th} = 10^{-5}/10^{-4} = 10^{-1} = 0.1$. So

$$10^9 \cdot (0.1)^{(d+1)/2} \leq 10^{-2} \implies (0.1)^{(d+1)/2} \leq 10^{-11}.$$

Since $(0.1)^n = 10^{-n}$, we need $(d+1)/2 \geq 11$, i.e.

$$d \geq 21, \quad R = \frac{d-1}{2} = 10,$$

where $d = 21$ satisfies the requirement with equality ($P_{fail} = 0.01$ exactly); a strict inequality requires $d \geq 22$.

Check: $(0.1)^{11} = 10^{-11}$; $\epsilon_L = 10^{-11}$ per cycle; over $T = 10^9$ cycles, $P_{fail} \approx T \epsilon_L = 10^9 \times 10^{-11} = 10^{-2} = 0.01$ (boundary case; any $d > 21$ gives margin).

4.3 Derivation 3: operating point required for a shallow code ($d = 11$)

For $d = 11$, exponent $(d+1)/2 = 6$. Requirement:

$$10^9 \left(\frac{\epsilon}{10^{-4}} \right)^6 \leq 10^{-2} \implies \left(\frac{\epsilon}{10^{-4}} \right)^6 \leq 10^{-11} \implies \frac{\epsilon}{10^{-4}} \leq 10^{-11/6}.$$

Compute $10^{-11/6}$: $-11/6 = -1.833333\dots$; $10^{-1.833333\dots} = e^{-1.833333\dots \times \ln 10} = e^{-1.833333 \times 2.302585} = e^{-4.221406}$. Since $e^{-4} = 0.018316$ and $e^{-0.221406} = 0.801345$, we get $e^{-4.221406} = 0.018316 \times 0.801345 = 0.014678$. Therefore

$$\epsilon \leq 0.014678 \times 10^{-4} = 1.4678 \times 10^{-6}.$$

Conclusion: a distance-11 implementation meets the reliability target only if the physical error rate is at most $\approx 1.47 \times 10^{-6}$, i.e., roughly 68 times below the hypothesized threshold ($10^{-4}/1.4678 \times 10^{-6} \approx 68.13$; check: $1/0.014678 = 68.13$).

4.4 Derivation 4: qubit overhead at $d = 21$ on $\mathcal{T}_3$

With $q = 3$, $R = 10$:

$$V_3(10) = 1 + 3 \frac{3^{10} - 1}{3 - 1} = 1 + 3 \frac{59049 - 1}{2} = 1 + 3 \times 29524 = 1 + 88572 = 88573.$$

Arithmetic check: $3^{10} = 59049$; $59049 - 1 = 59048$; $59048/2 = 29524$; $3 \times 29524 = 88572$; $88572 + 1 = 88573$.

So $N_{phys} = 88,573$ physical sites per logical qubit, i.e. 8.8573×10^4 .

4.5 Derivation 5: measurement bandwidth at $d = 21$

$$B = V_3(10) f_c = 88573 \times 10^6 \text{ s}^{-1} = 8.8573 \times 10^{10} \text{ s}^{-1}$$

per logical qubit. Arithmetic: $88573 \times 10^6 = 8.8573 \times 10^{10}$.

4.6 Derivation 6: overhead at the marginal operating point $\epsilon = 5 \times 10^{-5}$

Suppose the hardware achieves $\epsilon = 5 \times 10^{-5}$ (half the hypothesized threshold; stated assumption). Then $\epsilon/\epsilon_{th} = 0.5$ and

$$\epsilon_L = (0.5)^{(d+1)/2}.$$

Requirement $T \epsilon_L \leq 0.01$ gives $(0.5)^{(d+1)/2} \leq 10^{-11}$, i.e. $(d+1)/2 \geq \log_2(10^{11}) = 11 \log_2 10 = 11 \times 3.321928 = 36.541$. So $(d+1)/2 \geq 37$ (integer), $d \geq 73$, $R = 36$. Then

$$V_3(36) = 1 + 3 \frac{3^{36} - 1}{2}.$$

$3^{36} = (3^{18})^2 = 387,420,489^2 = 150,094,635,296,999,121$. Arithmetic check: $387,420,489^2 = (3.87420489 \times 10^8)^2 \approx 1.50095 \times 10^{17}$, consistent. Then

$$V_3(36) = 1 + 3 \frac{150094635296999121 - 1}{2} = 1 + 3 \times 75047317648499560 = 1 + 225141952945498680 = 225141952945498681.$$

Arithmetic check: $150094635296999121 - 1 = 150094635296999120$; $150094635296999120/2 = 75047317648499560$; $3 \times 75047317648499560 = 225141952945498680$; $+1 = 225141952945498681$.

So at the half-threshold operating point, $N_{phys} = 225141952945498681 \approx 2.2514 \times 10^{17}$ physical sites per logical qubit, and the measurement bandwidth is

$$B = V_3(36) f_c = 225141952945498681 \times 10^6 \text{ s}^{-1} \approx 2.2514195294549868 \times 10^{23} \text{ s}^{-1}.$$

Arithmetic check: $2.25141952945498681 \times 10^{17} \times 10^6 = 2.25141952945498681 \times 10^{23}$.

Conclusion: halving the margin to the threshold (from 10^{-1} to 5×10^{-1} in ϵ/ϵ_{th}) multiplies the required exponent $(d+1)/2$ from 11 to 37, and the required volume explodes from 8.8573×10^4 to $\approx 2.2514 \times 10^{17}$ sites — a factor of $225141952945498681/88573 \approx 2.5419 \times 10^{12}$. Operating close to the threshold is catastrophically expensive in this model; the resource accounting of [10] is therefore not a footnote but the dominant design constraint.

4.7 Derivation 7: consistency of the two conventions

Under Convention 2, after L RG levels the quadratic map iterates to $p_L = \kappa^{2^L-1} p_0^{2^L}$. The effective suppression exponent doubles each level, which is the same qualitative behavior as the power law $\epsilon_L \propto (\epsilon/\epsilon_{th})^{(d+1)/2}$ with d growing linearly in depth: both conventions express that logical failure is a high-order event in the physical error count. The conventions differ in how the order n of the correctable-error event scales with resources ($n = (d+1)/2$ linear in R under Convention 1; $n = 2^L$ exponential in depth under Convention 2 with constant κ). We therefore do not claim the conventions are equivalent; we use Convention 1 for all resource numbers and Convention 2 only for the threshold-fixed-point statement.

4.8 Derivation 8: threshold as unstable fixed point of the quadratic map

Fixed points of $p_{\ell+1} = \kappa p_\ell^2$ satisfy $p^* = \kappa(p^*)^2$, i.e. $p^*(\kappa p^* - 1) = 0$, giving $p_1^* = 0$ and $p_2^* = 1/\kappa$. For $p_0 < 1/\kappa$ the iteration $p_\ell = \kappa^{2^\ell-1} p_0^{2^\ell}$ converges to 0 (correctable); for $p_0 > 1/\kappa$ it diverges (uncorrectable). Identifying the physical error rate scale with $p_0 \sim \epsilon$ per (A1'), the threshold is

$$\epsilon_{th} = \frac{1}{\kappa}.$$

The hypothesis $\epsilon_{th} \geq 10^{-4}$ is therefore equivalent to

$$\kappa \leq \frac{1}{10^{-4}} = 10^4.$$

Arithmetic: $1/10^{-4} = 10^4$. **Conclusion:** under Convention 2 the 10^{-4} target holds if and only if the effective amplification constant satisfies $\kappa \leq 10^4$; measuring or bounding κ for the actual RG decoder of [9] is the single decisive experiment for this convention.

5. Results

All numbers below are derived in Section 4 under assumptions (A1)–(A7) (Convention 1) or (A1')–(A3') (Convention 2); none is a measurement.

- **R1 (Derivation 1).** At $\epsilon = \epsilon_{th}^{(0)} = 10^{-4}$ exactly, $\epsilon_L = A = 1$ per cycle for every d : zero protection. The threshold claim is meaningful only for ϵ strictly below 10^{-4} .
- **R2 (Derivation 2).** At $\epsilon = 10^{-5}$, the reliability target $P_{fail} \leq 0.01$ over $T = 10^9$ cycles requires $d \geq 21$ ($R = 10$).
- **R3 (Derivation 4).** At $d = 21$ on $\mathcal{T}_3$: $N_{phys} = V_3(10) = 88,573$ sites per logical qubit.
- **R4 (Derivation 5).** At $d = 21$, $f_c = 10^6 \text{ s}^{-1}$: $B = 8.8573 \times 10^{10} \text{ s}^{-1}$ measurement events per second per logical qubit.
- **R5 (Derivation 3).** A distance-11 code meets the same target only if $\epsilon \leq 1.4678 \times 10^{-6}$, i.e. a factor of ≈ 68.13 below the hypothesized threshold.
- **R6 (Derivation 6).** At $\epsilon = 5 \times 10^{-5}$ (half the hypothesized threshold), the target requires $d \geq 73$ ($R = 36$), $N_{phys} = V_3(36) = 225141952945498681 \approx 2.2514 \times 10^{17}$ sites, and $B \approx 2.2514195294549868 \times 10^{23} \text{ s}^{-1}$ per logical qubit.
- **R7 (Derivations 7–8).** Under Convention 2, $\epsilon_{th} = 1/\kappa$; the 10^{-4} target is equivalent to $\kappa \leq 10^4$.

6. Discussion

Limitations of the model. Every result is conditional on (A1)–(A7) and (A1')–(A3'). The most fragile assumptions are: (i) **independence (A1, A2')**. Real error channels on any physical substrate exhibit correlated and leakage-like errors; a quadratic map with constant κ and a clean power law both assume away correlations, and correlations typically *lower* the threshold. (ii) **Prefactor $A = 1$ (A2).** Any realistic architecture has $A > 1$, which pushes the required distance up by $\Delta d \approx 2 \ln A / |\ln(\epsilon/\epsilon_{th})|$; at $\epsilon/\epsilon_{th} = 0.1$, even $A = 10$ adds $\Delta d \approx 2 \times 2.302585/2.302585 = 2$, i.e. $d \geq 23$ and a larger volume. (iii) **Cycle rate (A6).** $f_c = 10^6 \text{ s}^{-1}$ is assumed, not derived; the bandwidth figures R4 and R6 scale linearly with f_c . (iv) **Decoder model.** We compute no decoding circuit or classical throughput; the workload characteristics that [3] and [2] identify (low data reuse, irregular graph access) suggest the classical decoding cost may rival the quantum measurement cost, but we have not modeled it. (v) **Convention 2 is heuristic.** The quadratic map is a leading-order recursion for bounded-weight uncorrectable patterns; if uncorrectable events at a node require errors on an unbounded number of children, the recursion order changes and $\epsilon_{th} = 1/\kappa$ fails.

Failure modes. The dominant quantitative failure mode is operating near the threshold: R6 shows that moving from $\epsilon/\epsilon_{th} = 0.1$ to 0.5 inflates the site count by a factor of $\approx 2.542 \times 10^{12}$. A second failure mode is underestimating κ : if the actual RG step amplifies with $\kappa > 10^4$, Convention 2 predicts no threshold at 10^{-4} at all.

What would falsify the claims. The conditional analysis is falsified as a *defense of the threshold hypothesis* if: (a) an explicit error model for the [9] construction yields $\epsilon_{th} < 10^{-4}$; (b) the prefactor A or correlated errors push the required operating point below what any hardware roadmap delivers; or (c) $\kappa > 10^4$ is derived for the actual RG decoder. Conversely, the analysis is falsified as a *criticism* if an explicit construction demonstrates $\epsilon_{th} \geq 10^{-4}$ with A and correlations accounted for, and a decoder whose classical bandwidth fits within the budget that [3], [7], and [8] suggest is achievable for graph-structured workloads.

Open questions. What is the actual value of κ for RG decoding on $\mathcal{T}_p$? How do correlations transform the recursion? Can decoding be distributed across subtrees in the manner suggested by the distributed-serving analogy of [1], and does the design-space perspective of [6] — evaluating the family $\{\mathcal{T}_p, d, \epsilon\}$ rather than one point — change the verdict? What lifecycle costs beyond steady-state operation, gestured at by the title of [12], dominate end to end? The supplied summaries of [4], [11], and [12] are too thin to support quantitative claims about these questions; we flag them as open.

7. Conclusion

We reconstructed, fully explicitly and conditionally, what a 10^{-4} threshold claim for holographic QEC as RG flow on Bruhat–Tits trees [9] must commit to. Under the stated model: exactly at $\epsilon = 10^{-4}$ there is no protection; at $\epsilon = 10^{-5}$ a distance-21 code suffices at 88,573 sites and $8.8573 \times 10^{10} \text{ s}^{-1}$ measurement events per logical qubit; a shallow $d = 11$ code needs $\epsilon \leq 1.4678 \times 10^{-6}$; and operating at half the threshold inflates the volume to $\approx 2.25 \times 10^{17}$ sites. In the complementary quadratic-RG convention, the threshold equals $1/\kappa$, so the target holds iff $\kappa \leq 10^4$. The threshold hypothesis is internally consistent but operationally demanding; its fate rests on measurable quantities — κ , A , and the correlation structure — that the current proposal does not supply.

References

- [1] Stochastic Model Based Proxy Servers Architecture for VoD to Achieve Reduced Client Waiting Time. arXiv:1002.1195v1. <https://arxiv.org/abs/1002.1195v1> [2] Architectural Implications of Graph Neural Networks. arXiv:2009.00804v2. <https://arxiv.org/abs/2009.00804v2> [3] Dalorex: A Data-Local Program Execution and Architecture for Memory-bound Applications. arXiv:2207.13219v4. <https://arxiv.org/abs/2207.13219v4> [4] Reasoning with Pixel-level Precision: QVLM Architecture and SquID Dataset for Quantitative Geospatial Analytics. arXiv:2601.13401v1. <https://arxiv.org/abs/2601.13401v1> [5] LSQCA: Resource-Efficient Load/Store Architecture for Limited-Scale Fault-Tolerant Quantum Computing. arXiv:2412.20486v2. <https://arxiv.org/abs/2412.20486v2> [6] Towards Assessing Spread in Sets of Software Architecture Designs. arXiv:2402.19171v1. <https://arxiv.org/abs/2402.19171v1> [7] A Novel Reconfigurable Architecture of a DSP Processor for Efficient Mapping of DSP Functions using Field Programmable DSP Arrays. arXiv:1306.0089v1. <https://arxiv.org/abs/1306.0089v1> [8] Insights into DeepSeek-V3: Scaling Challenges and Reflections on Hardware for AI Architectures. arXiv:2505.09343v2. <https://arxiv.org/abs/2505.09343v2> [9] QNFO: Holographic Quantum Error Correction as AdS/CFT Renormalization-Group Flow on Bruhat–Tits Trees: A Conditional Threshold Analysis at 10^{-4} [10] QNFO: The Resource Burden of Active Quantum Error Correction: A Scaling Analysis of Qubit, Bandwidth, and Energy Overheads [11] QNFO: Coherence-Assisted Error Correction for Deep Learning: A Reconciled Feasibility Analysis of a Hybrid Quantum-Classical Architecture [12] QNFO: Lifecycle of a Fault-Tolerant Quantum Computer

Appendix A. Divergence report

D1 (CONVERGENT with a documented convention choice): primary error model. One convention (Convention 1, adopted as primary) models the code by its distance $d = 2R + 1$ and the concatenation-style law $\epsilon_L = A(\epsilon/\epsilon_{th})^{(d+1)/2}$; a second convention (Convention 2, retained as complementary) models the RG step directly as $p_{\ell+1} = \kappa p_\ell^2$. The disagreement is one of modeling convention, not substance: Convention 1 makes resource accounting (volume, bandwidth) explicit, while Convention 2 makes the threshold a fixed point, $\epsilon_{th} = 1/\kappa$. The main text uses Convention 1 for all resource numbers and Convention 2 only for the fixed-point condition $\kappa \leq 10^4$ (Derivation 8); Derivation 7 records that the two are not quantitatively equivalent.

D2 (DIVERGENT, resolved): interpretation at $\epsilon = \epsilon_{th}$ exactly. One reading states that at $\epsilon = \epsilon_{th}$ the logical error per cycle is $A = 1$, i.e. zero protection (Derivation 1); the other states that at $p_0 = p^* = 1/\kappa$ the error neither shrinks nor grows under iteration (an unstable fixed point). Both statements are correct within their conventions; we adopt the stronger operational phrasing (zero protection) in the main text and note the fixed-point reading there, since a logical error rate of 1 per cycle and a non-converging iteration describe the same boundary behavior from different models.

Appendix B. Claim attribution

Claim	Substance	Source drafts	Status
C1	Threshold question framed by [9]; conditional analysis at 10^{-4}	A, B, C	CONVERGENT
C2	Resource-burden accounting frame from [10] (qubits, bandwidth, energy)	A, B	CONVERGENT
C3	Assumptions (A1)–(A7): independent errors, power law, $T = 10^9$, $P_{fail} \leq 0.01$, $f_c = 10^6 \text{ s}^{-1}$, $q = 3$	A, B, C	CONVERGENT
C4	Zero protection exactly at $\epsilon = 10^{-4}$	A, B, C	CONVERGENT
C5	$d \geq 21$ at $\epsilon = 10^{-5}$; $V_3(10) = 88,573$; $B = 8.8573 \times 10^{10} \text{ s}^{-1}$	A, B, C	CONVERGENT
C6	$d = 11$ requires $\epsilon \leq 1.4681 \times 10^{-6}$ (factor ≈ 68.1 below threshold)	A, B	CONVERGENT
C7	Derivation 6 at $\epsilon = 5 \times 10^{-5}$; $d \geq 73$, $V_3(36) = 225141952945498681$, $B \approx 2.2514 \times 10^{23} \text{ s}^{-1}$	A	SINGLE
C8	Convention 2: $p_{\ell+1} = \kappa p_\ell^2$, $\epsilon_{th} = 1/\kappa$, target iff $\kappa \leq 10^4$	B, C	CONVERGENT
C9	Interpretation at threshold: $\epsilon_L = 1$ vs unstable fixed point	A vs B	DIVERGENT (Appendix A, D2)
C10	Architectural analogies from [1]–[8], [11], [12] as co-design/memory/DSP/serving context	A, B, C	CONVERGENT